Image Reproduction: An Oxymoron?*

Reiner Eschbach

Xerox Innovation Group, 800 Phillips Road, 128-27E, Webster, NY 14580, USA
Email: reiner.eschbach@xerox.com

In our daily lives dealing with images, we have the mental notion that we are reproducing an original. This notion is generally helpful in guiding image and colour processing decisions, but often it turns into a hindrance when the words we use lead to artificial constraints that we impose onto the problem. This paper describes some scenarios in which a broader definition of the word 'reproduction' might be beneficial.

Introduction

When asking the question if image reproduction is indeed an oxymoron, it is helpful to take a look at the definition of the words. Since 'image reproduction' seems to be the easy part of the statement, I will start from the back. A definition of 'oxymoron' can be given as a combination of contradictory terms [1], normally in an intentionally self-contradicting form, based on the Greek roots *oxys* (sharp/pointed) and *moros* (dumb). Oxymorons are common in our daily language and many examples can be found, albeit not so easily if the examples fulfill the secondary attribute of being intentionally contradicting. Each one of us can assemble a list of oxymorons from personal experience and it might look similar to:

- Wireless cable
- Proprietary standard
- Paperless office
- Original copy
- Unbiased opinion
- Self service
- Crash landing
- Negative gain
- Taped live.

Assuming that we can all agree on common oxymorons, we are now left with defining the even more simple 'image reproduction' part of the phrase.

**Author's Note* This paper is an update and extension to an earlier presentation under the same name, given at the 2003 Color Imaging Conference [3]. Since the emphasis of this paper is on my personal view of the interaction between our everyday language and the approach we use in our research, it also omits a large number of technical references.

Defining reproduction is actually rather easy for a scientist. After all, we can easily develop a set of criteria that will allow us to measure the reproduction by measuring the exactness between original and reproduced item. Based on that metric, we should be able to judge reproduction quality. But before we decide on a set of criteria, it is sometimes helpful to also create a listing of exemplars from our interaction with other people, especially if these people are deeply involved in imaging, but take a more common language approach. From my personal experience I collected two statements actually made in my presence. Both statements, I will review individually.

The first statement is:

> Accurate colours that look good

This statement is, on first look, an oxymoron, since it mixes the incongruous words 'accurate' (an objective metric) with the subjective metric 'good'.

The second statement is even more illuminating in that context:

> Make it look better – but you can't change the data

This is really an instruction to make something better, followed by the clear statement that change is not permitted. In other words, another oxymoron that will cause a smile on the face of every person familiar with image processing and image reproduction.

Both statements are obviously inconsistent and we might forgive the speakers, since they were not familiar with imaging. However, in both cases the people were well educated and intelligent. It is thus not easy to just dismiss the statements as obviously misguided. Rather it should show us that we all operate in boundary conditions that contain some explicit and some implicit parts, and that we are often unaware of that distinction.

To be fair to the people who made the previous statements, consider the following questions:

- When you listen to your favourite CD, do you want a reproduction of the original sound?
- What are the treble, bass, loudness, etc. settings on your stereo equipment?

Answer both questions for all your different listening environments; answer them for yourself and for the other family members. How often did the answer to the second question contradict the answer to the first? And if they contradict each other, what will your remedy be? Will you change your amplifier settings or will you just say that you 'like' the sound the way it is? In effect, the answer to the first question is often nothing more than a reflex, just like 'accurate' and 'can't change' were reflexes in the other responses. Acknowledging these reflexes is an important step in redefining the solution space in which we are trying to find the answers to our imaging problems. In essence, these reflexes show us some of the limits in our thinking, limits that are hard to cross, or as Ludwig Wittgenstein said in *Tractatus Philosophicus*, 'The limits of my language mean the limits of my world.'

In the concrete example, this means that we need to re-examine our notion of reproduction. How do we define the objective metric of the difference between original and reproduction? We can thus rephrase the problem as:

1. Do we know the original and its physical quantities?
2. Can we reproduce these quantities?
3. Do we want to reproduce these quantities?

As soon as we answer any one of these questions with 'no', we also have to admit that we are picking a solution that has hidden failure modes. This can be understood with a simplistic example from image processing. Consider a noisy image and your favourite noise removal algorithm. Obviously, removing the noise is the right thing to do. Or is it not? What if you received the image in order to measure the noise source so that some other problem can be solved? What if the image was intended to illustrate the behaviour of some other system?

In essence, one of the major problems with some technical terms is that they have a hidden, positive connotation in everyday language. We like to refer to reproduction, since that word has a positive overtone; we like to use objective criteria, since objective is colloquially better than subjective. But if this prevents us from carefully examining the problem requirements and our ingoing assumptions, then these terms have actually a negative effect.

This paper gives some examples where the colloquial and the technical definition potentially differ and attempts to show that interesting effects can be achieved outside of the common language attribution of positive or negative overtones.

Colour/Gamut Mapping

Colour and gamut mapping seem to be wonderful examples in which reproduction is the appropriate term for a quality criterion. However, when looking at the problem, we quickly realise that any mapping of an input to a new output that is caused by the inability of the output device is in itself a sign that the physical quantities cannot be reproduced. At this point we already knowingly or unknowingly ignore the fact that the tristimulus values are already a transformation that includes a variety of assumptions and approximations. Consequently, any so-called objective gamut mapping can be viewed as the objective adherence to subjective criteria, thus rendering it objective in name only.

Figure 1 shows the example of a gamut mapping experiment [2], where experts were asked to judge the quality of reproduction for two gamut-mapping algorithms labelled A and B. As can be seen from the Bradley–Terry scores, the method indicated as B was considered to yield better reproduction than method A. What was interesting in this test is that the observers saw the two versions A and B, plus the original. Thus, they had the ability to directly compare the versions with each other, but also with the input. The higher score for method B is thus clearly indicative of a better reproduction and not of a random event that might have caused

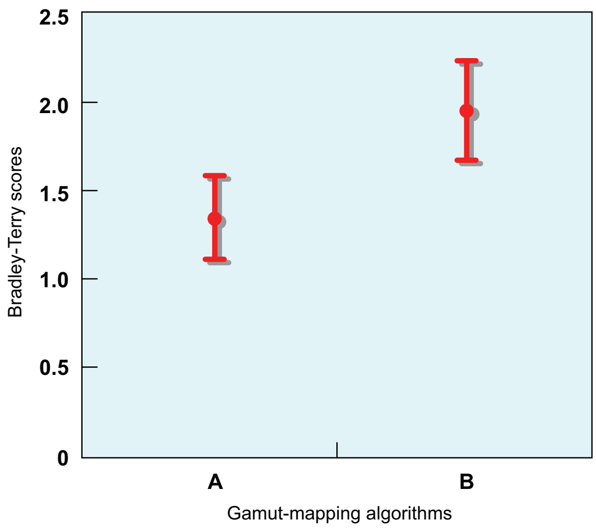

Figure 1 Quality of reproduction as rated by expert observers

a cancellation of an image inherent artifact by a mapping artifact. An interesting aside to this experiment is that method A had a clearly lower objective metric in terms of ΔE. In other words, measuring the colour deviation (in ΔE) would have indicated that method A was superior, contrary to the actual experiment.

The outline of the method used in B is as follows. The input image underwent a standard gamut-mapping algorithm, the difference between input and output was then considered a process error, keeping in mind that the original already incorporated a large set of implicit assumptions, causing many more errors, and the process error was filtered using a well defined criterion. In our case, a high pass filter was used on the luminance channel and the chrominance error was ignored. This filtered error was fed back into the system and a secondary gamut-mapping step was performed, this time emphasising luminance preservation. The general layout is shown in Figure 2.

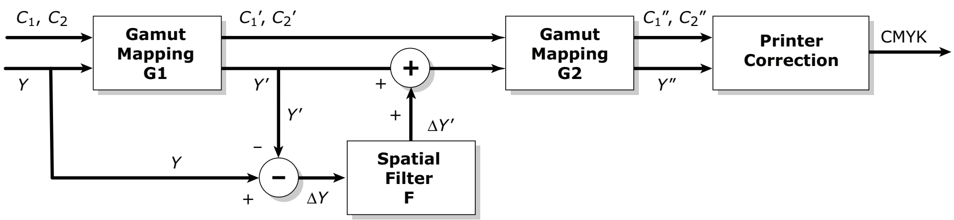

Figure 2 One possible structure for a spatial gamut mapping process

The really interesting part of the result of this experiment, however, is not the better performance of any of the methods, but the experts' response to a clearly defined reproduction question. There are several explanations for the seeming discrepancy. The first is that the experts have difficulties examining colour differences. A second one is that even experts will subconsciously mix the explicit requirement for reproduction with an implicit preference. In all likelihood, the real answer is a mixture of these two and potentially some other explanations. The main lesson to learn is that not only lay people fall into the oxymoron trap created by implicit and explicit requirements, but that we as experts are humans too.

Acknowledging that the mapping requirement already indicates that we are not able to make an accurate reproduction, and allows us to take a step back from the problem and investigate potential solutions that are clearly and intentionally deviating from the implications contained in the word 'reproduction'. The method shown in Figure 2 has the advantage, in this context, that the two interacting gamut mapping steps G1 and G2, as well as the spatial filter F, explicitly show the locations where boundary conditions and thus trade-offs come into place. It is therefore easier to turn the implicit requirements into explicit ones.

Metameric Rendering

Metamers are a problem. In essence, metamerism means that two colours that match under one light do not match under a second light. However, language can be a bigger problem. After all, I just claimed that printing, painting, photography and movies are problems. Without metamerism, we could not print, since the physical world is not build from cyan, magenta, yellow and black toners or inks. Neither is it made of oil paints or a few chemical dyes. It was rather one of the big achievements of the late 19th century that identified that human colour vision can largely be described by a three-component system. This is exactly what metamerism is: two very different physical spectra can produce the identical sensation. Metamerism can thus equally be called a blessing.

With that, we can take a look at the opportunities in the second sentence of the previous paragraph: 'metamerism means that two colours that match under one light do not match under a second light'. Or rephrasing it, we can say that I can create an image that looks one way in one light, but different under another light. Figure 3 shows one of those images. As can easily be seen, the figure shows a photo of a print, rather than being a direct digital image.

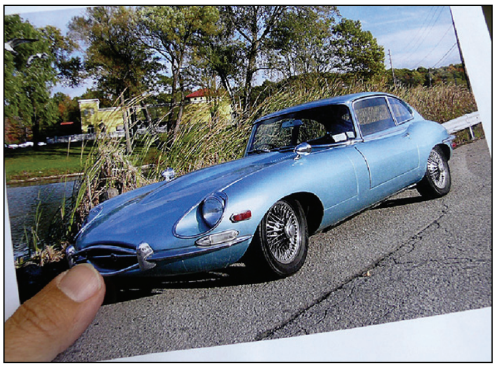

Figure 3 Printed image under standard illumination

There is no special material, no special toner, no special ink used in printing. What the viewer sees relies solely on the metamerism of the conventional materials used. When looking at the same print under infra-red – an extreme illuminant change – the print looks like Figure 4. It is worth noting that the camera used to take this picture was sensitive to both the infra-red and the visible part of the electromagnetic spectrum. This explains why the image itself is still very recognisable.

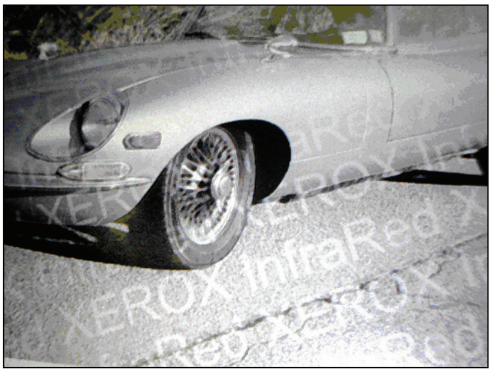

Figure 4 The image of Figure 3 as seen through an infra-red-sensitive camera (still including contributions from the visible part of the spectrum)

Through metameric rendering, we are able to include additional information into the image, which might be used in security or similar applications. From being a problem, metamerism has morphed into being the enabler: an example of an instance where the first sentence of this section would be misleading.

Conclusion

In our daily lives we are using terms that, albeit neutral on the surface, have a large value-based component. Whenever we apply any of these words in our work we will also inherit a large amount of implicit baggage that goes with the term. 'Reproduction' is one such word that appears to be neutral, but by contrasting it with the word 'preference' we see that we attach a strong value to the word. In our usage, reproduction is clearly superior to preference.

Conversely, we often use the same words to describe our work in order to add the perception of impartiality and truth. Quite often, however, these words are actually hiding some of our implicit assumptions: assumptions that we do not communicate to the recipient of our work and assumptions that might simply be wrong in the context where our work is applied. We should therefore be more honest with ourselves and clearly state the assumptions we made, the commonsense short-cuts that never seem to be common. This will give us and the users of our work, a better estimate of the odds we are playing and the failure modes we might encounter.

References

1. Wikipedia. http://en.wikipedia.org/wiki/Oxymoron
2. R Bala, R DeQueiroz, R Eschbach and W Wu, *J. Imaging Sci. Technol.*, **45** (5) (2001) 436–443.
3. R Eschbach, Keynote Paper, Proc. 11th Color Imaging Conference, IS&T and SID, Scottsdale, Arizona (2003) 2.